

NAMING CONVENTIONS

INTRODUCTION

At Microsoft we have been using a set of useful conventions for choosing identifiers in programs. You may be interested in these conventions because:

1. They can help in your programming task - no matter what language you use.
2. You may wish to read code written using the conventions.
3. You may desire to see the wider adoption of programming practices that promote the readability and portability of programs.

This monograph is intended to give you the flavor of the major ideas behind the conventions.

When confronted with the need for a new name in a program, a good programmer will generally consider the following factors to reach a decision:

1. Mnemonic value - so that the programmer can remember the name
2. Suggestive value - so that others can read the code
3. "Consistency" - this is often viewed as an aesthetic idea, yet it also has to do with the information efficiency of the program text. Roughly speaking, we want similar names for similar quantities.
4. Speed of the decision - we can't spend too much time pondering about the name of a single quantity, nor is there time for typing and editing extremely long variable names.

All in all, name selection can be a frustrating and time consuming subtask. Often a name which satisfies some of the above criteria will contradict the others. Maintaining consistency can be especially difficult.

ADVANTAGES OF THE CONVENTIONS

The following naming conventions provide a very convenient framework for generating names that satisfy the above criteria. The basic idea is to name all quantities by their types. This simple statement requires considerable elaboration (What is meant by "types"? What happens if "types" are not unique?) but once we can agree on the framework, the benefits readily follow. To wit:

1. The names will be mnemonic in a very specific sense: if someone remembers the type of a quantity or how it is constructed from other types, the name will be readily apparent.
2. The names will be suggestive as well: we will be able to map any name into the type of the quantity, hence obtaining information about the shape and the use of the quantity.
3. The names will be consistent because they will have been produced by the same rules.
4. The decision on the name will be mechanical, thus speedy.
5. Expressions in the program can be subjected to consistency checks that are very similar to the "dimension" checks in physics.

TYPE CALCULUS

As suggested above, the concept of "type" in this context is determined by the set of operations that can be applied to a quantity. The test for type equivalence is simple: could the same set of operations be meaningfully applied to the quantities in question? If so, the types are thought to be the same. If there are operations which apply to a quantity in exclusion of others, the type of that quantity is different. The concept of "operation" is considered quite generally here; "being the subscript of array A" or "being the second parameter of procedure Position" are operations on quantity x (and "A" or "Position" as well.) The point is that "integers" x and y are not of the same type if $Position(x,y)$ is legal but $Position(y,x)$ is nonsensical. Here we can also sense how the concepts of type and name merge: x is named so because it is an x-coordinate, and it seems that its type is also an x-coordinate. Most programmers would have probably named such a quantity x. In this instance, the conventions merely codify and clarify what has been widespread programming practice.

Note that the above definition of type (which, incidentally, is suggested by languages such as SIMULA and Smalltalk) is a superset of the more common definition which takes only the quantity's representation into account. Naturally, if the representations of x and y are different, there will exist some operations that could be applied to x but not y, or vice versa.

Let us not forget that we are talking about conventions which are to be used by humans for the benefit of humans. Capabilities or restrictions of the programming environment are not at issue here. The exact determination of what constitutes a "type" is not critical, either. If a quantity is misclassified, we have style problem, not a bug.

NAMING RULES

My thesis discusses in detail the following specific naming rules:

1. Quantities are named by their type possibly followed by a qualifier. A convenient (and legal) punctuation is recommended to separate the type and qualifier part of a name. (In C, we use a capital initial for the qualifier as in `rowFirst`: `row` is the type, `First` is the qualifier).
2. Qualifiers distinguish quantities that are of the same type and that exist within the same naming context. Note that contexts may include the whole system, a block, a procedure, or a data structure (for fields), depending on the programming environment. If one of the "standard qualifiers" is applicable, it should be used. Otherwise, the programmer can choose the qualifier. The choice should be simple to make, because the qualifier needs to be unique only within the type and within the scope - a set that is expected to be small in most cases. In rare instances more than one qualifier may appear in a name. Standard qualifiers and their associated semantics are listed below. An example is worthwhile: `rowLast` is a type `row` value; that is, the last element in an interval. The definition of "Last" states that the interval is "closed"; i.e., a loop through the interval should include `rowLast` as its last value.
3. Simple types are named by short tags that are chosen by the programmer. The recommendation that the tags be small is startling to many programmers. The essential reason for short tags is to make the implementation of rule 4 realistic. Other reasons are listed below.
4. Names of constructed types should be constructed from the names of the constituent types. A number of standard schemes for constructing pointer, array, and difference types exist. Other constructions may be defined as required. For example, the prefix `p` is used to construct pointers. `rowLast` is then the name of a particular pointer to a `row` type value that defines the end of a closed interval. The standard type constructions are also listed below.

In principle, the conventions can be enriched by new type construction schemes. However, the standard constructions proved to be sufficient in years of use. It is worth noting that the types for data structures are generally NOT constructed from the tags of their fields. First of all, constructions with over two components would be unwieldy. More importantly, the invariant property of data structures - the set of operations that they participate in - seems to be largely independent of the fields of the structure that determine only the representation. We all have had numerous experiences with changes in data structures that left the operations (but not the implementation of the operations) unchanged. Consequently, I recommend the use of a new tag for every new data structure. The tag with some punctuation (upper case initial or all upper case) should also be used as the structure name in the program. New tags should also be used if the constructions accumulate to the point where readability suffers.

In my experience, tags are more difficult to choose than qualifiers. When a new tag is needed, the first impulse is to use a short descriptive common generic English term as the type name. This is almost always a mistake. One should not preempt the most useful English phrases for the provincial purposes of any given version of a given program. Chances are that the same generic term could be equally applicable to many more types in the same program. How will we know which is the one with the pretty "logical" name, and which have the more arbitrary variants typically obtained by omitting various vowels or by other disfigurement? Also, in communicating with the programmer, how do we distinguish the generic use of the common term from the reserved technical usage? By inflection? In the long run, an acronym that is not an English word may work out the best for tags. Related types may then share some of the letters of the acronym. In speech, the acronym may be spelled out, or a pronounceable nickname may be used. When hearing the special names, the informed listener will know that the special technical meaning should be understood. Generic terms should remain free for generic use.

EXAMPLE

For example, a color graphics program may have a set of internal values that denote colors. What should one call the manifest value for the color red? The obvious choice (which is "wrong" here) is RED. The problem with RED is that it does not identify its type. Is it a label? A procedure that turns objects RED? Even if we know that it is a constant (because it is spelled all caps, for example) there might be several color-related types. Of which one is RED an instance? If I see a procedure defined as paint(color), may I call it RED as an argument? Has the word RED been used for any other purpose within the program? So we decide to find a tag for the color type and use the word Red as a qualifier.

Note that the obvious choice for the qualifier is in fact that the "correct" one! This is because the use of qualifiers are not hampered by any of the above difficulties. Qualifiers are not "exclusive" (or rather they are exclusive only within a smaller set) so we essentially need not take into account the possibility of other uses of the term "Red." The technical use of the term will be clear to everyone when the qualifier is paired up with an obviously technical type tag. Since qualifiers (usually) do not participate in type construction, there is no inherent reason why they would need to be especially short.

Conversely, the tag for the type of the color value should not be "color." Just consider all the other color related types that may appear in the graphics program (or in a future variant): hardware encoding of color; color map entry number; absolute pointer to color map entry; color values in alternate color mapping mode; hue-brightness-saturation triples; other color values in external interfaces: printers, plotters, interacting external software; etc. Furthermore, the tag will have to appear in names with constructed types and qualifiers.

A typical arbitrary choice could be "co" (pronounced see-oh). Or, if "co" was already taken, "cv", "cl", "kl", and so on. Note that the mnemonic value of the tags is just about average: not too bad, but not too good either. The conventions cannot help with creating names that are inherently mnemonic, instead they identify, compress, and contain those parts of the program that are truly individual, thus arbitrary. The lack of inherent meaning should be compensated by ample comments whenever a new tag is introduced. This is a reasonable suggestion since the number of basic tags remains very small even in a large system.

In conclusion, the name of our quantity would be "coRed", provided that the color type "co" is properly documented. The value of the name will show later in program segments such as:

```
if co == coRed then *mpcopx[coRed]+=dx ...
```

At a glance we can see that the variable co is compared with a quantity of its own kind; coRed is also used as a subscript to an array whose domain is of the correct type. Furthermore, as we will see, the color is mapped into a pointer to "x", which is de-referenced (by the * operator in this example) to yield an x type value, which is then incremented by a "delta x" type value. Such "dimensional analysis" does not guarantee that the program is completely free from bugs, but it does help to eliminate the most common kinds. It also lends a certain rhythm to the writing of the code: "Let's see, I have a co in hand and I need an x; do I have a mpcox? No, but there is a mpcopx that will give me a px; *px will get me the x...", and so on.

NAMING FOR "WRITABILITY"

A good yardstick for choosing a name is to try to imagine that there is an extraordinary reward for two programmers if they can independently come up with the same program text for the same problem. Both programmers know the reward, but cannot communicate otherwise. Such an experiment would be futile, of course, for any sizeable problem, but it is a neat goal. The reward in real life is that a program written by someone else, which is identical to what one's own program would have been, is extremely readable and modifiable. By the proper use of the conventions, the ideal can be approached very closely, give or take a relatively few tags and possibly some qualifiers. The leverage of the tags is enormous. If they are communicated, or are agreed on beforehand, or come from a common source, the goal becomes reachable and the reward may be reaped. This makes the documentation of the tags all the more important.

An example of such a consideration is the discretionary use of qualifiers in small scopes where a quantity's type is likely to be unique, for example in small procedures with a few parameters and locals or in data structures which typically have only a few fields. One might prefer to attach a qualifier even to a quantity with a unique type of way of explanation. While such redundancy cannot harm readability, it will hamper "writability", the ability for someone else to come up with the name without hesitation. As many textbooks point out, the "someone else" can be the same programmer sometime in the future revisiting the long forgotten code. Conclusion: don't use qualifiers when not needed, even if they seem valuable.

NAMING RULES FOR PROCEDURES

Unfortunately, the simple notion of qualified type tags does not work well for procedure names. Some procedures do not take parameters or do not return values. The scopes of procedure names tend to be large. The following set of special rules for procedures has worked quite satisfactorily:

1. Distinguish procedure names from other names by punctuation, for example by always starting with a capital letter (type tags of other quantities are in lower case). This alleviates the problem caused by the large scope.
2. Start the name with the tag of the value that is returned, if any.
3. Express the action of the procedure in one or two words, typically transitive verbs. The words should be punctuated for easy parsing by the reader (a common legal method of punctuation is the use of capital initials for every word).
4. Append the list of tags of some or all of the formal parameters if it seems appropriate to do so.

The last point is contrary to the earlier remarks on data structure naming. When the parameters to a procedure are changed, typically all uses of the procedure will have to be updated. There is an opportunity during the update to change the name as well - in fact the name change can serve as a useful check that all occurrences have been found. With data structures, the addition or change of a field will not have an effect on all uses of the changed structure type. Typically, if a procedure has only one or two parameters, the inclusion of the parameter tags will really simplify the choice of procedure name.

Some examples for procedure names:

InitSy: takes an sy as its argument and initializes it.

OpenFn: fn is the argument. The procedure will "open" the fn. No value is returned.

FcFromBnRn: returns the fc corresponding to the bn,m pair given. (The names cannot tell us what the types sy, fn, fc, etc., are.)

LIST OF STANDARD TYPE CONSTRUCTIONS

(X and Y stand for arbitrary tags. According to standard punctuation the actual tags are lower case.)

pX pointer to X

dX difference between two instances of type X. X + dX is of type X.

cX count of instances of type X.

mpXY an array of Y's indexed by X. Read as "map from X to Y."

rgX an array of X's. Read as "range X." The indices of the array are called:

iX index of the array rgX.

dnX (rare) an array indexed by type X. The elements of the array are called:

eX (rare) element of the array dnX.

grpX a group of X's stored one after another in storage. Used when the X elements are of variable size and standard array indexing would not apply. Elements of the group must be referenced by means other than direct indexing. A storage allocation zone, for example, is a grp of blocks.

bX relative offset to a type X. This is used for field displacements in a data structure with variable size fields. The offset may be given in terms of bytes or words, depending on the base pointer from which the offset is measured.

Where it matters, quantities named mp, rg, dn, or grp are actually pointers to the structures described above.

cbX size of instances of X in bytes

cwX size of instances of X in words

One obvious problem with the constructions is that they make the parsing of the types ambiguous. Is pfc a tag of its own or is it a pointer to an fc? Such questions (just as many others) can be answered only if one is familiar with the specific tags that are used in a program.

STANDARD QUALIFIERS

(The letter X stands for any type tag. Actual type tags are in lower case.)

XFirst the first element in an ordered set (interval) of X values.

XLast the last element in an ordered set of X values. XLast is the upper limit of a closed interval, hence the loop continuation condition should be: $x \leq x_{\text{Last}}$

XLim the strict upper limit of an ordered set of X values. Loop continuation should be: $x < x_{\text{Lim}}$

XMax strict upper limit for all X values (excepting Max, Mac, and Nil) for all other x: $x < x_{\text{Max}}$. If x values start with x=0, xMax is equal to the number of different x values. The allocated length of a dnx vector, for example, will be typically xMax.

XMac the Current (as opposed to constant or allocated) upper limit for all x values. If x values start with 0, xMac is the current number of X values. To iterate through a dnx array, for example:

for x=0 step 1 to xMac-1 do ... dnx[x] ...

or

for ix=0 step 1 to ixMac-1 do ... rgx[ix] ...

XNil a distinguished Nil value of type X. The value may or may not be 0 or -1.

XT temporary X. An easy way to qualify the second quantity of a given type in a scope.

SOME COMMON PRIMITIVE TYPES

f flag (boolean, logical). If qualifier is used, it should describe the true state of the flag. Exception: the constants fTrue and fFalse.

w word with arbitrary contents.

ch character, usually in ASCII text.

b byte, not necessarily holding a coded character, more akin to w. Distinguished from the b constructor by the capital letter of the qualifier in immediately following.

sz pointer to first character of a zero terminated string.

st pointer to a string. First byte is the count of characters cch.
h pp (in heap).

EXAMPLE

The following partial example of an actual symbol table routine illustrates the use of the conventions in a "real life" situation. The purpose of this example is not to make any claims about the code itself, but to show how the conventions can help us learn about the code. In fact, some of the names in this routine are standard.

```
1  #include "sy.h"
2  Extern int *rgwDic;
3  extern int bsyMac;
4  struct SY *PsySz(sz)
5  char sz[];
6  {
7      char *pch;
8      int cch;
9      struct SY *psy, *PsyCreate();
10     int *pbsy;
11     int cwSz;
12     unsigned wHash=0;
13     pch=sz;
14     while (*pch!=0)
15         wHash=(wHash<<5)+(wHash>>11)+*pch++;
16     cch=pch-sz;
17     pbsy = &rgbsyHash[(wHash&077777)%cwHash];
18     for (; *pbsy!=0; pbsy = &psy->bsyNext)
19     {
20         char *szSy;
21         szSy= (psy=(struct SY *)&rgwDic[*pbsy])->sz;
22         pch=sz;
23         while (*pch==*szSy++)
24             {
25                 if (*pch++==0)
26                     return (psy);
27             }
```

```

28          }
29      cwSz=0;
30      if (cch>=2)
31          cwSz=(cch-2/sizeof(int)+1);
32      *pbsy=(int *)PsyCreate(cwSY+cwSz)-rgwDic;
33      Zero((int *)psy, cwSY);
34      bltbyte(sz, psy->sz, cch+1);
35      return(psy);
36  }

```

The tag SY is the only product specific type in this routine. The definition of SY is found in the include file sy.h (fair enough). The type name itself is in all capitals, a common convention.

Line 2 says that there is an array of words, which is called Dic(tionary). Remember that since Dic is a qualifier, it is named traditionally.

Line 3 is the offset pointing beyond the last sy (see b constructor + Mac standard qualifier.) One has to guess at this time that this is used for allocating new sy's. The "base" of the offset would also have to be guessed to be rgwDic. Actually, the name grpsy would have been better instead of rgwDic, from this local perspective. In the real program, the rgwDic area is used for a number of different purposes, hence the "neutral" name.

4 is the procedure declaration. Procedure returns a pointer to an SY as the result. The parameter must be a zero terminated string.

7-12 declare quantities. The usages should be clear from the names. For example, cwSz is the number of words in some string (probably the argument), pbsy is a pointer to an offset of an sy (p constructor + b constructor). The only qualifier used here is in wHash - the hash code.

13 pch will be a pointer to the first character of sz.

16 cch is the count of characters (c constructor) ostensibly, in sz.

17 cwHash is the number of words in the hash table (I would have called it ibsyMax). In a way, the qualifier on rgbsyHash could be omitted, but it helps identifying the hash table in external contexts.

17-18 note the opportunities for dimensional checking:

pbsy = rgbsy[...] follows from pX = &rgX[...]

pbsy = psy->bsy Next follows from pX=&pY->X; or pX = &Y.X

so even the use of -> instead of . follows from local context. The p on the left hand side signals the need for the & on the right.

20 introduces a new sz, qualified to distinguish it from the argument. The qualifier, very appropriately, is the source of the datum, Sy.

23 given the use of szSy in this line, the name pchSy would have been a little more appropriate. No harm done, however.

29-31 This strange code has to do with the fact that the declaration of SY includes 2 bytes of sz, so that cwSz is really the number of words in the sz-2 bytes! This should deserve a comment or at least a qualifier M2 (minus 2) or the like. cwSY is the length of the SY structure in words. The all caps qualifier is not strictly standard, but it helps to associate the quantity with the declaration of SY, rather than with any random sy instance.

PsyCreate is a good procedure name; PsyCreateCw would have been even better. In line 32 we can also see an example of dimensional checking: while we have a psy inside the parenthesis, we need a bsy for the left side (*pbsy = bsy!) so we subtract the "base" of the bsy from the psy

$$bX + base = pX; \text{ hence: } bX = pX - base.$$

In closing, it is evident that the conventions participated in making the code more correct, easier to write and easier to read. Naming conventions cannot guarantee "good" code however, only the skill of the programmer can.

-- Charles Simonyi

6/24/83:800v